Lattice Mixing and Vanishing Trapdoors

A Framework for Fully Secure Short Signatures and more

Xavier Boyen

Universitas Leodiensis
Institut Montefiore
Liège, Belgium
xb@boyen.org

Abstract. We propose a framework for adaptive security from hard random lattices in the standard model. Our approach borrows from the recent Agrawal-Boneh-Boyen families of lattices, which can admit reliable and punctured trapdoors, respectively used in reality and in simulation. We extend this idea to make the simulation trapdoors cancel not for a specific target but on a non-negligible subset of the possible challenges. Conceptually, we build a compactly representable, large family of input-dependent "mixture" lattices, set up with trapdoors that "vanish" for a secret subset wherein we hope the attack occurs. Technically, we tweak the lattice structure to achieve "naturally nice" distributions for arbitrary choices of subset size. The framework is very general. Here we obtain fully secure signatures, and also IBE, that are compact, simple, and elegant.

1 Introduction

Lattices are currently enjoying renewed interest in cryptography, owing to a combination of mathematical elegance, implementation simplicity, provable security reductions, and, more recently, rather dramatic gains in efficiency that bring them closer to the familiar discrete-log and factoring-based approaches. Lattice-based crypto also offers the hope of withstanding quantum computers, against which both discrete-log and factoring-based approaches are known to be utterly defenseless. As a few examples of influential lattice-based cryptosystems and foundations, we mention [5, 6, 17, 18, 4, 23, 22, 19], among many others.

Still, by far the biggest barrier to the practical deployment of lattice-based cryptographic systems remains their space inefficiency, which may exceed by several orders of magnitude that of the mainstream. This is especially true for systems based on so-called "hard" random integer lattices, which have essentially no structure other than being periodic modulo the same modulus q along every coordinate axis. Hard lattices have the drawback of requiring voluminous representations, especially when compared to lattices with additional structure such as cyclic or ideal lattices. Being devoid of structure, however, hard lattices may harbor potentially tougher "hard problems" for a safer foundation for crypto.

A primary motivation for lattice cryptography being a hedge against the doomsday of mainstream assumptions, it seems worthwhile to endeavor to build cryptosystems as efficient and provably secure as we can from hard lattices.

1.1 Related Work

A number of advances toward efficient lattice-based signatures have recently been made. We mention a few that are most closely related to this work.

Lyubashevsky and Micciancio [16] gave an elegant near-linear-space one-time signature on cyclic lattices, that could be lifted via a standard tree construction into a many-time signature (in a stateful way, hence not hash-and-sign).

Gentry et al. [13] were the first to realize identity-based encryption from (hard) lattices, with a fully secure construction that also gave a very efficient signature as a by-product (albeit all of it only in the random-oracle model).

Cash et al. [11] and Peikert [21] then managed to remove the random oracle and add a hierarchy, using an elegant but bandwidth-intensive bit-by-bit scheme (also concurrently proposed by Agrawal and Boyen [3] *sans* hierarchy), reminiscent of Canetti et al. [10]. Additionally, Peikert [21] showed how the bit-by-bit framework could yield a simpler signature (but not IBE, as it was "salted"), using the recent prefix signature technique of Hohenberger and Waters [14].

Boneh et al. [9, 1] finally showed how to avoid the bit-by-bit IBE keys altogether, in favor of a compact and efficient all-at-once encoding, creating a selectively secure scheme reminiscent of Boneh and Boyen [8]. Though it did not natively give a secure signature (a costly generic conversion would be needed), we mention it because our framework turns it into a fully secure IBE and more.

Bandwidth Requirements The following table compares the space efficiency of those recent signature schemes (in hard lattices, unless indicated otherwise).

	S.I.S. strength β	Stand. model?	State-less?	|VerKey| # in $\mathbb{Z}_q$	|SigKey| # in $\mathbb{Z}$	|**Signature**| # in $\mathbb{Z}$
LM'08 [16]	✗ CYCL.LATTIC.	✓	✗ TREE	# in $\mathbb{Z}$: $\tilde{O}(\lambda)+\tilde{O}(\ell)$		$\tilde{O}(\ell)$
GPV'08 [13]	$\tilde{O}(n^{1.5})$	✗ R.O.M.	✓	$n\,m$	m^2	m
CHK'09 [11]	$\tilde{O}(\ell\,n^2)$	✓	✓	$2\,\ell^{2+\epsilon}\,n\,m$	m^2	$\ell^{2+\epsilon}\,m$
P'09 [21]	$\tilde{O}(\sqrt{\ell}\,n^{1.5})$	✓	✓	$2\,\ell\,n\,m$	m^2	$\ell\,m$
This work	$\tilde{O}(\ell\,n^{2.5})$	✓	✓	$\ell\,n\,m$	m^2	m

Parameter λ is the security level; ℓ the message bit-size; q the modulus; and m and n the lattice and constraint dimensions where $\lambda \approx n < m = \Theta(n \log q)$. Tabulated SIS strength is the approximation factor β incurred by the security reduction; and $|entity|$ the # of entries in $\mathbb{Z}$ and/or $\mathbb{Z}_q$ per *entity*, with up to $\lceil \log q \rceil \approx \lceil \log \beta \rceil$ bits per entry.

We remark that moderate differences of approximation parameter (β) should have limited practical impact compared to variations of number (#) of entries. Indeed, β is linked to the modulus of $\mathbb{Z}_q$ and the norm of entries in $\mathbb{Z}$; and varying their *magnitudes* by a factor $z = \ell^{k_1}\, n^{k_2} = \mathrm{poly}(q)$ only affects the information-theoretic bit sizes by a factor $1 + \log z / \log q = \Theta(1)$. By contrast, if we vary the *number* of entries by a factor z, the total bit sizes vary by a factor $\Theta(z)$. Note that $\forall \beta = \mathrm{poly}(n)$, there is an average-case β-SIS reduction [5] from worst-case SIVP with approximation factors $\gamma = \tilde{O}(\beta\sqrt{n})$, widely believed hard $\forall \gamma = \mathrm{poly}(n)$. The concrete parametric hardness of these assumptions is estimated in [12, 20].

1.2 Contribution

In this work, we propose a lattice-based encoding framework that generalizes the all-at-once encoding of Agrawal et al. [1]. The relationship of this work to the other one is akin to that linking Waters [24] to Boneh and Boyen [8] in pairing groups. Our goal is to build compact, practical, and "fully secure" signatures and identity-based encryption, from hard integer lattices in the standard model.

Here we focus on signatures. Our main construction is a stateless "hash-and-sign" fully secure signature, i.e., existentially unforgeable under chosen-message attacks, that is about as short as [13]. Our main result is a standard-model security reduction for it and related schemes (from the classic average-case SIS problem, itself reducible from worst-case SIVP and other hard problems [5, 23]).

As a bonus, our framework yields a clean "unsalted" construction that extends effortlessly from signature to identity-based private-key extraction. The two are indeed closely related, except that certain tricks used to make signatures secure are incompatible with IBE, such as black-box randomized hashes whose "nonces" would be inaccessible to a non-interactive encrypting party. Our framework does not have this problem, and has already been used to make the IBE scheme of [1] fully secure with little loss of efficiency (see the full version of [1] for details).

1.3 Highlights

Technically, we obtain our compact signature by "mixing" together, in a message-dependent manner, a number of public-key matrices in order to induce in a deterministic way a large family of hard lattices. A signature is a short non-zero vector in the appropriate lattice. For proving adaptive security, we arrange the lattice melange in such a way that a signing trapdoor, i.e., a short lattice basis, is always available for every possible input in the real scheme. In the simulation, faulty trapdoors will be made to vanish through spurious cancellations for a certain, suitably sized set of "challengeable inputs", unknown to the adversary.

A crucial and novel feature of our framework is to ensure that the challengeable inputs are well spread out over the entire input space, regardless of the selected size of the challengeable set. This ensures that, regardless of the actions of the adversary, the simulation will unfold with a significant and more or less invariant probability of success. This simulation robustness property is unusual and key to achieving an efficient security reduction.

Earlier schemes, also based on this principle of small but non-negligible challengeable input sets, generally did not have the luxury of uniform distributions over custom domains; they had to provision complex mechanisms to compensate for the non-uniformity of certain events in function of the adversary's actions. The Waters [24] scheme, for example, contains such a mechanism, prompted by the non-existence of distributions of non-negligible equal weights over exponentially sized groups as used in pairing-based cryptography.

With lattices, by contrast, the possibility to work with smaller moduli gives us an extra handle on the construction of "nice" distributions for a very wide range of challengeable input set sizes. As a result, we obtain security reductions that are simpler, tighter, and more efficient.

2 Lattice Notions

Here we gather a number of useful notions and results from the literature.

We denote by $\|\mathsf{A}\|$ or $\|\mathbf{a}\|$ the ℓ_2-norm of a matrix A or vector $\mathbf{a}$. We denote by $\tilde{\mathsf{A}}$ the Gram-Schmidt ordered orthogonalization of A, and its ℓ_2-norm by $\|\tilde{\mathsf{A}}\|$.

2.1 Random Integer Lattices

Definition 1. Let a basis $\mathsf{B} = \left[\ \mathbf{b}_1 \ \middle| \ \ldots \ \middle| \ \mathbf{b}_m \ \right] \in \mathbb{R}^{m \times m}$ be an $m \times m$ matrix with linearly independent columns $\mathbf{b}_1, \ldots, \mathbf{b}_m \in \mathbb{R}^m$. The lattice Λ generated by the basis B and its dual Λ^* are defined as (both are m-dimensional),

$$\Lambda = \mathcal{L}(\mathsf{B}) = \left\{ \ \mathbf{y} \in \mathbb{R}^m \quad \text{s.t.} \quad \exists \mathbf{s} \in \mathbb{Z}^m \ , \quad \mathbf{y} = \mathsf{B}\,\mathbf{s} = \sum_{i=1}^{m} \mathsf{s}_i\,\mathbf{b}_i \ \right\}$$

$$\Lambda^* = \left\{ \ \mathbf{z} \in \mathbb{R}^m \quad \text{s.t.} \quad \forall \mathbf{y} \in \Lambda \ , \quad \mathbf{z}^T\,\mathbf{y} = \langle \mathbf{z}, \mathbf{y} \rangle \ \in \mathbb{Z} \ \right\}$$

Definition 2. For a positive integer q (later a prime) and a matrix $\mathsf{A} \in \mathbb{Z}_q^{n \times m}$, define two m-dimensional full-rank integer lattices:

$$\Lambda^\perp(\mathsf{A}) = \left\{ \ \mathbf{e} \in \mathbb{Z}^m \quad \text{s.t.} \quad \mathsf{A}\,\mathbf{e} = 0 \pmod{q} \ \right\}$$

$$\Lambda(A) = \left\{ \ \mathbf{y} \in \mathbb{Z}^m \quad \text{s.t.} \quad \exists \mathbf{s} \in \mathbb{Z}^n \ , \quad \mathsf{A}^T\,\mathbf{s} = \mathbf{y} \pmod{q} \ \right\}$$

These are dual when properly scaled, as $\Lambda^\perp(\mathsf{A}) = q\,\Lambda(\mathsf{A})^*$ and $\Lambda(\mathsf{A}) = q\,\Lambda^\perp(\mathsf{A})^*$.

2.2 Bases and Trapdoors

A fundamental result in the geometry of numbers is that every lattice Λ has a basis; e.g., see [15]. Implicit in its proof, is the well known fact that any full-rank set $\mathsf{S}_A \subset \Lambda$ can be converted into a basis T_A for Λ with no greater orthogonalized norm $\|\tilde{\mathsf{T}}_A\| \le \|\tilde{\mathsf{S}}_A\|$.

Fact 3. *For a set $\mathsf{X} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m\}$ of lattice vectors, let $\tilde{\mathsf{X}} = \{\tilde{\boldsymbol{x}}_1, \ldots, \tilde{\boldsymbol{x}}_m\}$ be its Gram-Schmidt ordered orthogonalization. There is a deterministic polynomial-time algorithm that, on input an arbitrary basis of an m-dimensional lattice Λ and a full-rank set $\mathsf{S} = \{\boldsymbol{s}_1, \ldots, \boldsymbol{s}_m\} \subset \Lambda$ of lattice vectors, returns a basis $\mathsf{T} = \{\boldsymbol{t}_1, \ldots, \boldsymbol{t}_m\}$ of Λ such that $\|\tilde{\boldsymbol{t}}_i\| \le \|\tilde{\boldsymbol{s}}_i\|$ for all $i = 1, \ldots, m$.*

Ajtai [6] shows how to sample a uniform matrix $\mathsf{A} \in \mathbb{Z}_q^{n \times m}$ with an associated full-rank set $\mathsf{S}_A \subset \Lambda^\perp(\mathsf{A})$ of low-norm vectors orthogonal to A modulo q. Tightness was later improved by Alwen and Peikert [7].

Proposition 4 ([7])**.** For any $\delta_0 > 0$, there is a probabilistic polynomial-time algorithm that, on input a security parameter 1^λ, an odd prime $q = \mathrm{poly}(\lambda)$, and two integers $n = \Theta(\lambda)$ and $m \ge (5 + 3\,\delta_0)\,n \log q$, outputs a statistically $(m\,q^{-\delta_0\,n/2})$-close to uniform matrix $\mathsf{A} \in \mathbb{Z}_q^{n \times m}$ and a basis $\mathsf{T}_A \subset \Lambda^\perp(A)$ such that with overwhelming probability $\|\mathsf{T}_A\| \le O(n \log q)$ and $\|\tilde{\mathsf{T}}_A\| \le O(\sqrt{n \log q})$.

For the purpose of this paper, we take $\delta_0 = 1/3$, assume $L = \tilde{\Omega}(\sqrt{m})$, and summarize the foregoing as follows.

Fact 5. *There is a probabilistic polynomial-time algorithm that, on input a security parameter 1^λ, an odd prime $q = \mathrm{poly}(\lambda)$, and two integers $n = \Theta(\lambda)$ and $m \geq 6\, n \log q$, outputs a matrix $\mathsf{A} \in \mathbb{Z}_q^{n \times m}$ statistically close to uniform, and a basis T_A for $\Lambda^\perp(\mathsf{A})$ with overwhelming probability such that $\|\tilde{\mathsf{T}_A}\| \leq \tilde{\Theta}(\sqrt{m}) \leq L$.*

2.3 Discrete Gaussians

Given a basis for an integer random lattice, we recall how to sample random lattice points from a discrete Gaussian distribution whose minimum "width" is function of the norm of the lattice basis. We follow the works of [23, 4, 19, 13].

Definition 6. Let $m \in \mathbb{Z}_{>0}$ be a positive integer and $\Lambda \subset \mathbb{R}^m$ an m-dimensional lattice. For any vector $\mathbf{c} \in \mathbb{R}^m$ and any positive parameter $\sigma \in \mathbb{R}_{>0}$, we define:

$\rho_{\sigma,\mathbf{c}}(\mathbf{x}) = \exp\left(-\pi \frac{\|\mathbf{x}-\mathbf{c}\|^2}{\sigma^2}\right)$: a Gaussian-shaped function on $\mathbb{R}^m$ with center $\mathbf{c}$ and parameter σ, (For $x \in \mathbb{R}$, $\rho_{\sigma,c}(x) \propto \mathcal{N}_{\frac{\sigma}{\sqrt{2\pi}},0}(x)$, the normal probability density of variance $\frac{\sigma^2}{2\,\pi}$ and mean 0.)

$\rho_{\sigma,\mathbf{c}}(\Lambda) = \sum_{\mathbf{x}\in\Lambda} \rho_{\sigma,\mathbf{c}}(\mathbf{x})$: the (always converging) discrete integral of $\rho_{\sigma,\mathbf{c}}$ over the lattice Λ,

$\mathcal{D}_{\Lambda,\sigma,\mathbf{c}}$: the discrete Gaussian distribution over Λ with center $\mathbf{c}$ and parameter σ,

$$\forall y \in \Lambda \quad , \quad \mathcal{D}_{\Lambda,\sigma,\mathbf{c}}(y) = \rho_{\sigma,\mathbf{c}}(y) / \rho_{\sigma,\mathbf{c}}(\Lambda)$$

For notational convenience, origin-centered $\rho_{\sigma,\mathbf{0}}$ and $\mathcal{D}_{\Lambda,\sigma,\mathbf{0}}$ are abbreviated as ρ_σ and $\mathcal{D}_{\Lambda,\sigma}$.

Gentry et al. [13] show that, given a basis B for a lattice Λ, one can efficiently sample points in Λ with discrete Gaussian distribution for sufficiently large values of σ.

Proposition 7 ([13])**.** There exists a probabilistic polynomial-time algorithm that, on input an arbitrary basis B of an m-dimensional full-rank lattice $\Lambda = \mathcal{L}(\mathsf{B})$, a parameter $\sigma \geq \|\tilde{\mathsf{B}}\|\, \omega(\sqrt{\log m})$, and a center $\mathbf{c} \in \mathbb{R}^m$, outputs a sample from a distribution that is statistically close to $\mathcal{D}_{\Lambda,\sigma,\mathbf{c}}$.

For concreteness, we will refer to the algorithm of Proposition 7 as follows:

SampleGaussian$(\mathsf{B}, \sigma, \mathbf{c})$**:** On input a basis B for a lattice $\Lambda \subset \mathbb{R}^m$, a positive real parameter $\sigma \geq \|\tilde{\mathsf{B}}\|\, \omega(\sqrt{\log m})$, and a center vector $\mathbf{c} \in \mathbb{R}^m$, it outputs a fresh random lattice vector $\mathbf{x} \in \Lambda$ drawn from a distribution statistically close to $\mathcal{D}_{\Lambda,\sigma,\mathbf{c}}$.

2.4 Smoothing Parameter

We recall the notion of smoothing parameter of a lattice which lower-bounds the "density" of points on a lattice across all directions, and how this relates to discrete Gaussian sampling on the lattice.

Micciancio and Regev [19] define the smoothing parameter of a lattice as follows.

Definition 8 ([19]). For any m-dimensional lattice Λ and any positive real $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest real $\eta > 0$ such that $\rho_{1/\eta}(\Lambda^* \setminus \{0\}) \leq \epsilon$.

Micciancio and Regev [19] show that large deviations from lattice points vanish exponentially.

Proposition 9 ([19]). For any lattice Λ of integer dimension m, any point $\mathbf{c}$, and any two reals $\epsilon \in (0,1)$ and $\eta \geq \eta_\epsilon(\Lambda)$,

$$\Pr\left\{ \mathbf{x} \sim \mathcal{D}_{\Lambda,\eta,\mathbf{c}} \;:\; \|\mathbf{x} - \mathbf{c}\| > \sqrt{m}\,\eta \right\} \;\leq\; \frac{1+\epsilon}{1-\epsilon}\,2^{-m}$$

Peikert and Rosen [22] show that the Gaussian function itself vanishes away from any point.

Proposition 10 ([22]). For any lattice Λ of integer dimension m, any center $\mathbf{c} \in \mathbb{R}^m$, any two reals $\epsilon \in (0,1)$ and $\eta \geq 2\,\eta_\epsilon(\Lambda)$, and any lattice point $\mathbf{x} \in \operatorname{span}(\Lambda)$,

$$\mathcal{D}_{\Lambda,\eta,\mathbf{c}}(x) \;\leq\; \frac{1+\epsilon}{1-\epsilon}\,2^{-m}$$

2.5 Statistical Mixing

We recall some useful statistical mixing properties relating to the reduction of an integer vector modulo a lattice to yield a syndrome.

Ajtai [5] then Regev [23] show that binary combinations of enough vectors alsmost always span the space.

Proposition 11 ([23]). Let $m \geq 2\,n \log q$. Then for all except at most some q^{-n} fraction of matrices $\mathsf{A} \in \mathbb{Z}_q^{n\times m}$, the subset sums of the columns of A generate $\mathbb{Z}_q^n$. In other words, for every syndrome $\mathbf{u} \in \mathbb{Z}^n$ there exists a binary vector $\mathbf{e} \in \{0,1\}^m$ such that $\mathsf{A}\,\mathbf{e} = \mathbf{u} \pmod q$.

Gentry et al. [13] show that short Gaussian combinations of any spanning vector set yields uniformity.

Proposition 12 ([13]). Assume the columns of $\mathsf{A} \in \mathbb{Z}_q^{n\times m}$ generate $\mathbb{Z}_q^n$, and let $\epsilon \in (0,1)$ and $\eta \geq \eta_\epsilon(\Lambda^\perp(\mathsf{A}))$. Then for $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^m,\eta}$ the distribution of the syndrome $\mathbf{u} = \mathsf{A}\,\mathbf{e} \bmod q$ is within statistical distance $2\,\epsilon$ of uniform over $\mathbb{Z}_q^n$.

Furthermore, fix $\mathbf{u} \in \mathbb{Z}_q^n$ and let $\mathbf{c} \in \mathbb{Z}^m$ be an arbitrary solution to $\mathsf{A}\,\mathbf{c} = \mathbf{u} \pmod q$. Then the conditional distribution of $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^m,\eta}$ given $\mathsf{A}\,\mathbf{e} = \mathbf{u} \pmod q$ is exactly $\mathbf{c} + \mathcal{D}_{\Lambda^\perp(\mathsf{A}),\eta,-\mathbf{c}}$.

Gentry et al. [13] then show that for random A the lattice $\Lambda(\mathsf{A})$ has large minimal distance in ℓ_∞ and thus that $\Lambda^\perp(\mathsf{A})$ has small smoothing parameter.

Proposition 13 ([13])**.** Let q be a prime and n and m be two integers satisfying $m \geq 2\,n\,\log q$. Then, for all but at most some q^{-n} fraction of matrices $\mathsf{A} \in \mathbb{Z}_q^{n\times m}$, it holds that $\lambda_1^\infty(\Lambda(\mathsf{A})) \geq q/4$. Also, for any such A and any $\omega(\sqrt{\log m})$ function, there is a negligible function $\epsilon(m)$ such that the smoothing parameter $\eta_\epsilon(\Lambda^\perp(\mathsf{A})) \leq \omega(\sqrt{\log m})$.

Combining the previous propositions, Gentry et al. [13] summarize the results as follows.

Fact 14. *Fix a prime q and two integers n and m satisfying $m \geq 2\,n\,\log q$. For all but at most $2\,q^{-n}$ of matrices $\mathsf{A} \in \mathbb{Z}_q^{n\times m}$ and for any Gaussian parameter $\eta \geq \omega(\sqrt{\log m})$, on input $\boldsymbol{e} \sim \mathcal{D}_{\mathbb{Z}^m,\eta}$ the distribution of the syndrome $\boldsymbol{u} = \mathsf{A}\,\boldsymbol{e} \bmod q$ is statistically close to uniform over $\mathbb{Z}^n$.*

2.6 Preimage Sampling

We recall the notion of preimage-samplable functions (PSF) defined in [13], which is based on the combination of a trapdoor construction for integer lattices and an efficient discrete Gaussian sampling algorithm.

Let a uniform matrix $\mathsf{A} \in \mathbb{Z}_q^{n\times m}$ and a low-norm basis T_A for the lattice $\Lambda^\perp(\mathsf{A})$. Used in the discrete Gaussian sampling algorithm, the short basis T_A can act as a trapdoor for finding small non-zero solutions $\mathbf{e} \in \mathbb{Z}^m$ of the equation $\mathsf{A}^T\,\mathbf{e} = 0 \pmod q$ or more generally $\mathsf{A}^T\,\mathbf{e} = \mathbf{u} \pmod q$ for any $\mathbf{u} \in \mathbb{Z}_q^n$. This leads to the notion of preimage-samplable functions [13].

We give the following definition of preimage-samplable function, following [13]:

Definition 15. Let λ, q, n, m, and L be as in Fact 5. Let $\sigma \geq L\,\omega(\sqrt{\log m})$ be some Gaussian parameter. A preimage-samplable function family is a collection of maps $f_A : \mathbb{D}_{\mathbb{Z}^m,\sigma} \to \mathbb{Z}_q^n$ from $\mathbb{D}_{\mathbb{Z}^m,\sigma} = \{\mathbf{e} \in \mathbb{Z}^m : \|\mathbf{e}\| \leq \sqrt{m}\,\sigma\} \subseteq \mathbb{Z}^m$ into $\mathbb{Z}_q^n$, and specified by the following four algorithms:

$\mathsf{TrapGen}(1^\lambda)$: On input 1^λ, it uses the algorithm of Fact 5 to obtain a pair $(\mathsf{A}, \mathsf{T}_A)$, where $\mathsf{A} \in \mathbb{Z}_q^{n\times m}$ is statistically close to uniform and $\mathsf{T}_A \subset \lambda^\perp(\mathsf{A})$ is a short basis with $\|\tilde{\mathsf{T}}\| \leq L$. The public function parameters are (A, q). The preimage-sampling trapdoor is the basis T_A.

$\mathsf{EvalFun}(\mathsf{A}, q, \mathbf{e})$: On input function parameters (A, q) and an input point $\mathbf{e} \in \mathbb{D}_{\mathbb{Z}^m,\sigma}$, it outputs the image $f_A(\mathbf{e}) = \mathsf{A}\,\mathbf{e} \bmod q$ in $\mathbb{Z}_q^n$. (The output is undefined on large input $\mathbf{e} \in \mathbb{Z}^m \setminus \mathbb{D}_{\mathbb{Z}^m,\sigma}$.)

$\mathsf{SampleDom}(1^{(m)}, \sigma)$: On input the $m \times m$ identity matrix $1^{(m)}$ and a Gaussian parameter σ, it outputs $\mathbf{e} \leftarrow \mathsf{SampleGaussian}(1^{(m)}, \sigma, \mathbf{0})$, i.e., outputs an element $\mathbf{e} \in \mathbb{Z}^m$ such that $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^m,\sigma}$. The input matrix $1^{(m)}$ conveys the dimension m and its columns give a basis for Gaussian sampling in the lattice $\mathbb{Z}^m$. By Proposition 10, with overwhelming probability $\mathbf{e} \in \mathbb{D}_{\mathbb{Z}^m,\sigma}$.

SamplePre$(\mathsf{A}, q, T_A, \sigma, \mathbf{u})$**:** On input function parameters A and q and a trapdoor T_A, a Gaussian parameter σ as above, and a target image $\mathbf{u} \in \mathbb{Z}_q^n$, it samples a preimage $\mathbf{e} \in \mathbb{D}_{\mathbb{Z}^m,\sigma}$ from the distribution $\mathcal{D}_{\mathbb{Z}^m,\sigma}$ conditioned on the event that $\mathsf{A}\,\mathbf{e} = \mathbf{u} \pmod q$. To do this, it solves for an arbitrary solution $\mathbf{c} \in \mathbb{Z}^m$ in the linear system $\mathsf{A}\,\mathbf{c} = \mathbf{u} \pmod q$; it then samples $\mathbf{d} \leftarrow \mathsf{SampleGaussian}(\mathsf{T}_A, \sigma, -\mathbf{c}) \sim \mathcal{D}_{\Lambda^\perp(\mathsf{A}),\sigma,-\mathbf{c}}$ and outputs $\mathbf{e} = \mathbf{c} + \mathbf{d}$ in $\mathbb{Z}^m$. By Proposition 10, with overwhelming probability $\mathbf{e} \in \mathbb{D}_{\mathbb{Z}^m,\sigma}$.

The construction is correct and efficient by Proposition 12; see [13] for details.

2.7 Elementary Delegation

There are several ways to delegate a short basis for $\Lambda^\perp(\mathsf{A})$ into one for $\Lambda^\perp([\mathsf{A}|\mathsf{B}])$. If there is no one-wayness requirement on the delegation process, then Peikert [21] describes a very effective elementary deterministic way to do this.

Proposition 16 ([21])**.** Take any matrix $\mathsf{A} \in \mathbb{Z}_q^{n \times m_1}$ such that the columns of A span the group $\mathbb{Z}_q^n$. Let an arbitrary $\mathsf{B} \in \mathbb{Z}_q^{n \times m_2}$, and define $\mathsf{F} = [\mathsf{A}|\mathsf{B}]$. There exists a polynomial-time deterministic algorithm that, given A, B, and an arbitrary basis T_A for $\Lambda^\perp(\mathsf{A})$, outputs a basis T_F for $\Lambda^\perp(\mathsf{F})$ while preserving the Gram-Schmidt norm of the basis (i.e., such that $\|\tilde{\mathsf{T}}_F\| = \|\tilde{\mathsf{T}}_A\|$).

2.8 Hardness Assumption

The following lattice problem was first suggested to be hard on average by Ajtai [5] and formally defined by Micciancio and Regev [19].

Definition 17. The Small Integer Solution (SIS) problem in L_2-norm is: given an integer q, a matrix $\mathsf{A} \in \mathbb{Z}_q^{n \times m}$, and a real β, find a non-zero integer vector $\mathbf{e} \in \mathbb{Z}^m$ such that $\mathsf{A}\,\mathbf{e} = \mathbf{0} \pmod q$ and $\|\mathbf{e}\|_2 \leq \beta$. The average-case (q, n, m, β)-SIS problem is defined similarly, where A is uniformly random.

This problem was shown to be as hard as certain worst-case lattice problems, first by Ajtai [5], then by Micciancio and Regev [19], and Gentry et al. [13].

Proposition 18 ([13])**.** For any poly-bounded m, any $\beta = \mathrm{poly}(n)$ and for any prime $q \geq \beta \cdot \omega(\sqrt{n \log n})$, the average-case (q, n, m, β)-SIS problems is as hard as approximating the Shortest Independent Vector Problem (SIVP), among others, in the worst case to within certain $\gamma = \beta \cdot \tilde{O}(\sqrt{n})$ factors.

2.9 More Useful Facts

Lemma 19. *Let* $\mathsf{B}_0 \in \mathbb{Z}_q^{n \times m}$. *Let* H *be a scalar* $\mathsf{h} \in \mathbb{Z}_q$ *or a matrix* $\mathsf{H} \in \mathbb{Z}_q^{n \times n}$. *Suppose that* H *is invertible modulo* q *(i.e.,* $|H| \neq 0 \pmod q$ *when* q *is prime). Then, the two preimage-samplable functions* $(\mathsf{B}_0)(\cdot) \bmod q$ *and* $(H\,\mathsf{B}_0)(\cdot) \bmod q$ *from* $\mathbb{Z}^m$ *into* $\mathbb{Z}_q^n$ *admit exactly the same trapdoors* $\mathsf{T}_{B_0} \subset \mathbb{Z}^m$.

Proof. For all $\mathbf{e} \in \mathbb{Z}^m$ we have $\mathsf{B}_0\,\mathbf{e} = \mathbf{0} \pmod q$ if and only if $H\,\mathsf{B}_0\,\mathbf{e} = \mathbf{0} \pmod q$, hence the two lattices $\Lambda^\perp(\mathsf{B}_0)$ and $\Lambda^\perp(H\,\mathsf{B}_0)$ are the same. Thus, $\mathsf{T}_{B_0} \subset \Lambda^\perp(\mathsf{B}_0) \Leftrightarrow \mathsf{T}_{B_0} \subset \Lambda^\perp(H\,\mathsf{B}_0)$. $\square$

3 General Simulation Framework

We now describe the core scheme. At a high level, we achieve short signatures with full adaptive security by providing a relatively large number of public-key matrices, which are then "mixed through" together in a message-dependent manner — as opposed to merely juxtaposed as in the constructions of [3, 11, 21]. In the simulation, the public-key matrices will hide a trapdoor component that has a non-negligible probability of vanishing in the mix for certain unpredictable choices of messages: on those messages the simulator will be unable to answer signature queries, but will be able instead to exploit an existential forgery.

Our key-mixing technique is at some level reminiscent of Waters' scheme [24] in bilinear groups, but with a number of crucial differences. The farther-reaching difference is that in the lattice setting we can exploit the smaller groups and their richer structure to create a (much) more efficient "mixing" effect than in the large cyclic groups of the discrete-log setting. Another difference concerns randomization, which in a lattice setting tends to be rather more involved than in discrete-log settings; our approach is based on the method of randomization by a low-norm matrix from [1], with the small added contribution to show that it can be done in a way that supports the mixing effect that we need.

3.1 Two-Sided Trapdoors

To facilitate the description of the scheme and its proof, we first construct a preimage-samplable function of a special form that will be able to sample short preimages from the same distribution, using either one of two types of trapdoors: "firm" trapdoors will be used in the real scheme, and will never fail to work; "fickle" trapdoors will be used in the simulation, and will be fragile by design.

Lattices with dual trapdoors were first introduced in [9, 1]. Here, we seek to let the matrix R, below, be generated as a mixture of certain low-norm matrices. All the algorithms in this subsection are adapted from § 4 of [1].

Definition 20. Consider an algorithm $\mathsf{TwoSideGen}(1^\lambda)$ that outputs two random matrices $\mathsf{A} \in \mathbb{Z}_q^{n\times m}$ and $\mathsf{R} \in Z^{m\times m}$, where A is uniform and R has some distribution $\mathcal{R}$. Let $\mathsf{B} \in \mathbb{Z}_q^{n\times m}$ be an independent third matrix. Write $\mathsf{A}\,\mathsf{R}$ as shorthand for $(\mathsf{A}\,\mathsf{R} \bmod q) \in \mathbb{Z}_q^{n\times m}$, and define,

$$\mathsf{F} \;=\; \big[\mathsf{A} \;\big|\; \mathsf{A}\,\mathsf{R} + \mathsf{B}\big] \quad \in\; \mathbb{Z}_q^{n\times 2m}$$

We say that the pair (F, q) defines the public parameters of a *two-sided function*.

The following lemmas show that a two-sided function (F, q) is a preimage-samplable function given a trapdoor for either A or B, provided that A and R are drawn from suitable distributions.

Lemma 21. *For any parameter $\eta \geq \omega(\sqrt{\log m})$, there exists an efficiently samplable distribution $\mathcal{R}_\eta$ over $\mathbb{Z}^{m\times m}$, such that with overwhelming probability $\mathsf{R} = \sum_{i=1}^{\ell}$ for independent $\mathsf{R}_i \sim \mathcal{R}_\eta$ has norm $\|\mathsf{R}\| \leq \sqrt{m}\,\eta$, and such that for $(\mathsf{A}, \mathsf{R}) \sim \mathcal{U}_{\mathbb{Z}_q^{n\times m}} \times \mathcal{R}$ and fixed $\mathsf{B} \in \mathbb{Z}_q^{n\times m}$ the matrix $\mathsf{F} = [\mathsf{A}|\mathsf{A}\,\mathsf{R} + \mathsf{B}] \in \mathbb{Z}_q^{n\times 2m}$ is statistically close to uniform.*

Proof. According to Fact 14, it suffices to pick the columns of R independently wiht $\sim \mathcal{D}_{\mathbb{Z}^m,\eta}$. $\square$

Lemma 22 ("Firm" trapdoor)**.** *Let L and σ be as in Definition 15 and $\mathcal{R}_\eta$ as in Lemma 21. If $[\mathsf{A}|\mathsf{B}] \sim \mathcal{U}_{\mathbb{Z}_q^{n\times 2m}}$ and $\mathsf{T}_A \subset \Lambda^\perp(\mathsf{A})$ of norm $\|\tilde{\mathsf{T}}_A\| \leq L$, then the pair $\big(\mathsf{F} = [\mathsf{A}|\mathsf{B}],\ q\big)$ is a preimage-samplable function in the sense of Definition 15.*

Proof. Per Lemma 21, F is statistically close to uniform in $\mathbb{Z}_q^{n\times 2m}$, thus F has the right distribution. It remains to show how to perform public and trapdoor sampling.

$\mathsf{SampleDom}$. To sample short vectors $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{2m},\sigma}$ in the domain of F, one proceeds exactly as in the GPV scheme, i.e., by executing $\mathsf{SampleDom}(1^{(2m)},\sigma)$ which does not require any trapdoor.

$\mathsf{SamplePre}$. For preimage sampling, we show how to sample a short preimage $\mathbf{e} \in \mathbb{Z}^{2m}$ of any $\mathbf{u} \in \mathbb{Z}_q^n$ with conditional distribution $\mathcal{D}_{\mathbb{Z}^{2m},\sigma} \mid \mathsf{F}\,\mathbf{e} = \mathbf{u} \pmod q$. Since a random $\mathsf{A} \in \mathbb{Z}_q^{n\times m}$ will almost always span all of $\mathbb{Z}_q^n$, we can use the deterministic delegation mechanism of Proposition 16 to obtain a basis T_F for F with short Gram-Schmidt norm $\|\tilde{\mathsf{T}}_F\| \leq L$. Having such a trapdoor T_F for F, we invoke $\mathsf{SamplePre}(\mathsf{F}, q, \mathsf{T}_F, \sigma, \mathbf{u})$ to obtain a short random preimage $\mathbf{e}$. $\square$

Lemma 23 ("Fickle" trapdoor)**.** *Let L be as in Definition 15, η as in Lemma 21, and $\sigma = L'\,\omega(\sqrt{\log m})$ where $L' = 2\,\eta\,\sigma'\,\sqrt{m}$ and where $\sigma' \geq L\,\sqrt{\ell\, m}\,\omega(\sqrt{\log m})$. Fix a matrix $\mathsf{B} \in \mathbb{Z}_q^{n\times m}$ with a short basis T_B of orthogonalized norm $\|\tilde{\mathsf{T}}_B\| \leq L$. For (A, R) such that $[\mathsf{A}|\mathsf{A}\,\mathsf{R}] \sim \mathcal{U}_{\mathbb{Z}_q^{n\times 2m}}$ and $\|\mathsf{R}\| \leq \eta\,\sqrt{\ell\, m}$, the pair $\big(\mathsf{F} = [\mathsf{A}|\mathsf{A}\,\mathsf{R} + \mathsf{B}],\ q\big)$ is a preimage-samplable function in the sense of Definition 15.*

In this lemma, we allow $\|\mathsf{R}\| \leq \eta\,\sqrt{\ell\, m}$, where the factor $\sqrt{\ell}$ will account for the fact that in the simulation the matrix $\mathsf{R} = \mathsf{R}_{\mathsf{msg}} = \sum_{i=1}^{\ell} \pm \mathsf{R}_i$ for independent R_i of norm $\|\mathsf{R}_i\| \leq \eta\,\sqrt{m}$ and coefficients ± 1 function of the message msg.

Proof. Per Lemma 21, F is statistically close to uniform in $\mathbb{Z}_q^{n\times 2m}$, thus F has the right distribution. We need to show how to perform public and trapdoor sampling.

$\mathsf{SampleDom}$. Sampling short vectors $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{2m},\sigma}$ is done without any trapdoor by invoking $\mathsf{SampleDom}(1^{(2m)},\sigma)$, as in the previous lemma.

$\mathsf{SamplePre}$. For preimage sampling, we need to show, given any input $\mathbf{u} \in \mathbb{Z}_q^n$, how to sample a short preimage $\mathbf{e} \in \mathbb{Z}^{2m}$ of $\mathbf{u}$ with conditional distribution $\mathcal{D}_{\mathbb{Z}^{2m},\sigma} \mid \mathsf{F}\,\mathbf{e} = \mathbf{u} \pmod q$. We do this in three steps:

1. We build a full-rank set $\mathsf{S}_F \subset \Lambda^\perp(\mathsf{F})$ such that $\|\tilde{\mathsf{S}}_F\| \leq 2\,\eta\,L\,\sqrt{\ell}\,m\,\omega(\sqrt{\log m})$. This is done by independently sampling short vectors $\mathbf{e}_i \in \Lambda^\perp(\mathsf{F})$ until a linearly independent set of $2\,m$ such vectors is found. To sample one short vector $\mathbf{e} \in \Lambda^\perp(\mathsf{F})$ given the trapdoor T_B, we compute $\mathbf{d}_1 \leftarrow \mathsf{SampleDom}(1^{(m)}, (\eta\,\sqrt{\ell} - 1)\,\sigma')$ and $\mathbf{d}_2 \leftarrow \mathsf{SamplePre}(\mathsf{B}, q, \mathsf{T}_B, \sigma', -\mathsf{A}\,\mathbf{d}_1)$, and define,

$$\mathbf{d} \;=\; \begin{bmatrix}\mathbf{d}_1\\ \mathbf{d}_2\end{bmatrix} \;\in\; \mathbb{Z}^{2m} \qquad\qquad \mathbf{e} \;=\; \begin{bmatrix}\mathbf{d}_1 - \mathsf{R}\,\mathbf{d}_2\\ \mathbf{d}_2\end{bmatrix} \;\in\; \mathbb{Z}^{2m}$$

Observe that $\mathbf{e}$ is a fixed invertible linear function of $\mathbf{d}$, and that $\mathbf{d}$ is discrete Gaussian by construction. A result of Regev [23] shows that, with overwhelming probability, at most $4\,m^2$ samples will be needed to get $2\,m$ linearly independent vectors "$\mathbf{d}$", and therefore also $2\,m$ linearly independent vectors "$\mathbf{e}$". For each $\mathbf{e}$, we have $\mathsf{F}\,\mathbf{e} = \mathsf{A}\,(\mathbf{d}_1 - \mathsf{R}\,\mathbf{d}_2) + (\mathsf{A}\,\mathsf{R} + \mathsf{B})\,\mathbf{d}_2 = \mathsf{A}\,\mathbf{d}_1 + \mathsf{B}\,\mathbf{d}_2 = \mathsf{A}\,\mathbf{d}_1 - \mathsf{A}\,\mathbf{d}_1 = \mathbf{0} \in \mathbb{Z}_q^n$, hence $\mathbf{e} \in \Lambda^\perp(\mathsf{F})$. We have also $\|\mathbf{e}\| \leq (\eta\,\sqrt{\ell} - 1)\,\sigma'\,\sqrt{m} + \eta\,\sigma'\,m\,\sqrt{\ell} + \sigma'\,\sqrt{m} \leq 2\,\eta\,\sigma'\,m\,\sqrt{\ell}$. Thus by assembling $2\,m$ linearly independent such vectors "$\mathbf{e}$", we obtain a full-rank set $\mathsf{S}_F \subset \Lambda^\perp(\mathsf{F})$ of orthogonalized norm $\|\tilde{\mathsf{S}}_F\| \leq 2\,\eta\,\sigma'\,m\,\sqrt{\ell}$.

2. We convert the short set S_F into an equally short basis T_F, i.e., such that $\|\tilde{\mathsf{T}}_F\| \leq \|\tilde{\mathsf{S}}_F\|$. We can do this efficiently using the algorithm of Fact 3, starting from an arbitrary basis for $\Lambda^\perp(\mathsf{F})$, itself easy to construct by linear algebra.

3. We use the newly constructed basis T_F to sample a short preimage $\mathbf{e}$ of the given target $\mathbf{u} \in \mathbb{Z}_q^n$, using $\mathbf{e} \leftarrow \mathsf{SamplePre}(\mathsf{F}, q, \mathsf{T}_F, \sigma, \mathbf{u})$. Notice that the Gaussian parameter $\sigma \geq \|\tilde{\mathsf{T}}_F\|\,\omega(\sqrt{\log m})$, so the algorithm $\mathsf{SamplePre}$ can be applied with the stated parameters, and hence $\mathbf{e}$ sampled in this manner will have conditional distribution $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{2m},\sigma} \mid \mathsf{F}\,\mathbf{e} = \mathbf{u} \pmod q$. □

Remark 24. Agrawal et al. [2] show the sampling overhead is only a factor ≤ 2, hence in Step 1 we need to sample at most $4\,m$ vectors "$\mathbf{e}$" on expectation.

We also mention that a lower-norm fickle trapdoor may be obtained by using the Alwen-Peikert delegation method as in Lemma 22 instead of the repeated sampling as above. We revisit this issue in the full paper.

The point of the two-sided preimage-samplable function is that in the actual scheme we use the "firm" preimage mechanism with an always-available trapdoor T_A, whereas in the simulation we use the "fickle" preimage mechanism T_B for a matrix $\mathsf{B} = \mathsf{h}_{\mathsf{msg}}\,\mathsf{B}_0$ that sometimes vanishes.

3.2 Main Signature Scheme

The following is our core construction of a fully secure short signature. It is very simple and already achieves most of the compactness benefits while illustrating the framework. In the full paper, we show how to squeeze out some additional factor from the signature size, albeit at the cost of a more complex system.

From now on, a message msg is an ℓ-bit string $\big(\mathsf{msg}[1], \ldots, \mathsf{msg}[\ell]\big) \in \{0,1\}^\ell$ indexed from 1 to ℓ, augmented with a 0-th dummy extra bit set to $\mathsf{msg}[0] = 0$. This will let us easily include a constant term of index 0 in various summations.

KeyGen(1^λ)**:** On input a security parameter λ in unary, do these steps:

1. Draw an n-by-m matrix $\mathsf{A}_0 \in \mathbb{Z}_q^{n\times m}$ with a short basis $\mathsf{T}_{A_0} \subset \Lambda^\perp(\mathsf{A}_0)$.
   – Do so by invoking $\mathsf{TrapGen}(1^\lambda)$, resulting in T_{A_0} such that $\|\tilde{\mathsf{T}}_{A_0}\| \leq L$.
2. Draw $\ell + 1$ independent n-by-m-matrices $\mathsf{C}_0, \ldots, \mathsf{C}_\ell \in \mathbb{Z}_q^{n\times m}$.
3. Output the signing and verification keys,

$$\mathsf{SK} = \big(\mathsf{T}_{A_0}\big) \in \mathbb{Z}^{m\times m} \qquad \mathsf{VK} = \big(\mathsf{A}_0,\ \mathsf{C}_0, \ldots, \mathsf{C}_\ell\big) \in (\mathbb{Z}_q^{n\times m})^{\ell+2}$$

Sign$(\mathsf{SK}, \mathsf{msg})$**:** On input a signing key SK and a message $\mathsf{msg} \in \{0\} \times \{0,1\}^\ell$:

1. Define the n-by-m-matrix $\mathsf{C}_{\mathsf{msg}} = \sum_{i=0}^{\ell} (-1)^{\mathsf{msg}[i]} \, \mathsf{C}_i$.
2. Define the message-dependent matrix $\mathsf{F}_{\mathsf{msg}} = \left[\mathsf{A}_0 \;\middle|\; \mathsf{C}_{\mathsf{msg}}\right] \in \mathbb{Z}_q^{n \times 2m}$.
3. Sample a short non-zero random point $\mathbf{d} \in \Lambda^\perp(\mathsf{F}_{\mathsf{msg}})$, using $\mathsf{SK} = \mathsf{T}_{A_0}$.
– Do so by sampling $\mathbf{d} \sim \mathcal{D}_{\mathbb{Z}^{2m},\sigma} \mid \mathsf{F}_{\mathsf{msg}}\, \mathbf{d} = 0$, using Lemma 22.
4. Output the digital signature,

$$\mathsf{sig}_{\mathsf{msg}} \;=\; \big(\mathbf{d}\big) \quad \in \; \mathbb{Z}^{2m}$$

Verify$(\mathsf{VK}, \mathsf{msg}, \mathsf{sig}_{\mathsf{msg}})$**:** On input a verification key VK, a message msg, and a signature $\mathsf{sig}_{\mathsf{msg}}$:

1. Check that the message msg is well formed in $\{0\} \times \{0,1\}^\ell$.
2. Check that the signature $\mathsf{sig}_{\mathsf{msg}}$ is a small but non-zero vector.
– Do so by verifying that $\mathsf{sig}_{\mathsf{msg}} = \mathbf{d} \in \mathbb{Z}^{2m}$ and $0 < \|\mathbf{d}\| \le \sqrt{2\,m} \cdot \sigma$.
3. Check that $\mathsf{sig}_{\mathsf{msg}}$ is a point on the “mixed” lattice specified by msg.
– Do so by verifying that

$$\left[\mathsf{A}_0 \;\middle|\; \sum_{i=0}^{\ell} (-1)^{\mathsf{msg}[i]} \, \mathsf{C}_i\right] \mathbf{d} \;=\; \mathbf{0} \pmod{q}$$

4. If all the verifications pass, accept the signature; otherwise, reject.

3.3 Security Reduction

It is easy to see by inspection that the signature scheme is consistent with overwhelming probability.

The next theorem reduces the SIS problem to the existential forgery of our signature. The proof involves a moderate polynomial SIS parameter β. The expression of β arises in Lemma 26, but otherwise “passes through” the reduction. In § 3.4, we revisit the question of the lattice parameters in greater detail.

Theorem 25. *For a prime modulus $q = q(\lambda)$, if there is a probabilistic algorithm $\mathcal{A}$ that outputs an existential signature forgery, with probability ϵ, in time τ, and making $Q \le q/2$ adaptive chosen-message queries, then there is a probabilistic algorithm $\mathcal{B}$ that solves the (q, n, m, β)-SIS problem in time $\tau' \approx \tau$ and with probability $\epsilon' \ge \epsilon/(3\,q)$, for some polynomial function $\beta = \mathrm{poly}(\lambda)$.*

Proof. Suppose that there exists such a forger $\mathcal{A}$. We construct a solver $\mathcal{B}$ that simulates an attack environment and uses the forgery to create its solution. The various operations performed by $\mathcal{B}$ are the following.

Invocation. $\mathcal{B}$ is invoked on a random instance of the (q, n, m, β)-SIS problem, and is asked to return an admissible solution.

– Supplied: an n-by-m-matrix $\mathsf{A}_0 \in \mathbb{Z}_q^{n \times m}$ from the uniform distribution.
– Requested: any $\mathbf{e}_0 \in \mathbb{Z}^m$ such that $\mathsf{A}_0\, \mathbf{e}_0 = 0 \pmod{q}$ and $0 \ne \|\mathbf{e}_0\| \le \beta$.

Setup. $\mathcal{B}$ gives to the adversary $\mathcal{A}$ a simulated verification key constructed as follows:

1. Pick a random matrix $\mathsf{B}_0 \in \mathbb{Z}_q^{n\times m}$ with a short basis $\mathsf{T}_{B_0} \subset \Lambda^\perp(\mathsf{B}_0)$.
   – Do so by invoking $\mathsf{TrapGen}(1^\lambda)$, resulting in T_{B_0} such that $\|\tilde{\mathsf{T}}_{B_0}\| \leq L$.
2. Pick $\ell+1$ short random square m-by-m-matrices $\mathsf{R}_0, \ldots, \mathsf{R}_\ell \in \mathbb{Z}^{m\times m}$.
   – Do so by independently sampling the columns of the $\mathsf{R}_i \sim \mathcal{D}_{\mathbb{Z}^m,\eta}$.
3. Pick ℓ uniformly random scalars $\mathsf{h}_1, \ldots, \mathsf{h}_\ell \in \mathbb{Z}_q$ and fix $\mathsf{h}_0 = 1 \in \mathbb{Z}_q$.
4. Output the verification key $\mathsf{VK} = \big(\ \mathsf{A}_0,\quad \mathsf{C}_0 = (\mathsf{A}_0\,\mathsf{R}_0 + \mathsf{h}_0\,\mathsf{B}_0) \bmod q,$ $\mathsf{C}_1 = (\mathsf{A}_0\,\mathsf{R}_1 + \mathsf{h}_1\,\mathsf{B}_0) \bmod q,\ \ \ldots,\ \ \mathsf{C}_\ell = (\mathsf{A}_0\,\mathsf{R}_\ell + \mathsf{h}_\ell\,\mathsf{B}_0) \bmod q\ \big)$.

Queries. $\mathcal{B}$ answers adaptive signature queries from $\mathcal{A}$ on any message msg as follows:

1. Compute the matrix $\mathsf{R}_{\mathsf{msg}} = \sum_{i=0}^{\ell} (-1)^{\mathsf{msg}[i]}\,\mathsf{R}_i$.
2. Compute the scalar $\mathsf{h}_{\mathsf{msg}} = \sum_{i=0}^{\ell} (-1)^{\mathsf{msg}[i]}\,\mathsf{h}_i$.
3. If $\mathsf{h}_{\mathsf{msg}} = 0 \pmod q$, abort the simulation.
4. Compute the matrix $\mathsf{F}_{\mathsf{msg}} = \big[\mathsf{A}_0 \mid \mathsf{A}_0\,\mathsf{R}_{\mathsf{msg}} + \mathsf{h}_{\mathsf{msg}}\,\mathsf{B}_0\big] \in \mathbb{Z}_q^{n\times 2m}$.
5. Find a short random $\mathbf{d} \in \Lambda^\perp(\mathsf{F}_{\mathsf{msg}}) \subset \mathbb{Z}^{2m}$, using the trapdoor T_{B_0}.
   – Do so by sampling $\mathbf{d} \sim \mathcal{D}_{\mathbb{Z}^{2m},\sigma}$ given $\mathsf{F}_{\mathsf{msg}}\,\mathbf{d} = 0$, using the procedure of Lemma 23, using T_{B_0} as short basis for $\Lambda^\perp(\mathsf{h}_{\mathsf{msg}}\,\mathsf{B}_0)$ per Lemma 19.
6. Output the digital signature $\mathsf{sig}_{\mathsf{msg}} = \mathbf{d} \in \mathbb{Z}^{2m}$.

Forgery. $\mathcal{B}$ receives from $\mathcal{A}$ a forged signature $\mathbf{d}^*$ on a new (unqueried) message msg^*, and does:

1. Compute the matrix $\mathsf{R}^* = \sum_{i=0}^{\ell} (-1)^{\mathsf{msg}^*[i]}\,\mathsf{R}_i$.
2. Compute the scalar $\mathsf{h}^* = \sum_{i=0}^{\ell} (-1)^{\mathsf{msg}^*[i]}\,\mathsf{h}_i$.
3. If $\mathsf{h}^* \neq 0 \pmod q$, abort the simulation.
4. Separate $\mathbf{d}^{*T}$ into $\big[\mathbf{d}_1^{*T} \mid \mathbf{d}_2^{*T}\big]$.
5. Return $\mathbf{e}_0 = \mathbf{d}_1^* + \mathsf{R}^*\,\mathbf{d}_2^* \in \mathbb{Z}^m$ as solution to $\mathsf{A}_0\,\mathbf{e}_0 = \mathbf{0} \pmod q$.

Lemma 26 shows that the answer $\mathbf{e}_0$ will be with small and non-zero with good probability, and thus a valid (q, n, m, β)-SIS solution for the stated approximation β. (An instantiation of β is given in § 3.4.)

Outcome. The reduction is valid provided that $\mathcal{B}$ can complete the simulation (without aborting) with a substantial probability that is independent of the view of $\mathcal{A}$ and the choices it makes. The completion probability for $\mathcal{B}$ against an arbitrary strategy for $\mathcal{A}$ is quantified in Lemma 27.

It follows from the bounds of Lemmas 26 and 27, under the assumption that $Q \leq q/2$, that if $\mathcal{A}$ existentially forges a signature with probability ϵ, then $\mathcal{B}$ solves the SIS instance with probability,

$$\epsilon' \ \geq\ \pi_0\,\big(1 - q^{-1}\,Q\big)\,q^{-1}\,\epsilon\ \geq\ \pi_0\,\epsilon/2\,q\ \geq\ \epsilon/3\,q \qquad \text{for}\ \ \pi_0 \geq 2/3$$

With the stated lemmas, this concludes the security reduction. □

Lemma 26. *Given a valid forgery* $\big[\boldsymbol{d}_1^{*T} \mid \boldsymbol{d}_2^{*T}\big]$ *from* $\mathcal{A}$ *on some* msg^* *such that* $\mathsf{h}_{\mathsf{msg}^*} = 0 \pmod q$*, the vector* $\boldsymbol{e}_0 = \boldsymbol{d}_1^* + \mathsf{R}_{\mathsf{msg}^*}\,\boldsymbol{d}_2^* \in \mathbb{Z}^m$ *is with high probability* $\pi_0 = \Theta(1) \geq 2/3$ *a short non-zero preimage of* 0 *under* A_0*, namely,* $\boldsymbol{e}_0 \in \Lambda^\perp(\mathsf{A}_0)$ *and* $0 \neq \|\boldsymbol{e}_0\| \leq \beta$ *for some polynomial function* $\beta = \mathrm{poly}(\ell, n, m) = \mathrm{poly}(\lambda)$.

Sketch. Let $\mathsf{h}^* = \mathsf{h}_{\mathsf{msg}^*}$ and $\mathsf{R}^* = \mathsf{R}_{\mathsf{msg}^*}$. Let $\mathsf{C}^* = \mathsf{C}_{\mathsf{msg}^*} = \sum_{i=0}^{\ell} (-1)^{\mathsf{msg}^*[i]}\, \mathsf{C}_i$.

First, when $\mathsf{h}^* = 0$, we have $\mathsf{C}^* = \mathsf{A}_0\, \mathsf{R}^* + \mathsf{h}^*\, \mathsf{B}_0 = \mathsf{A}_0\, \mathsf{R}^*$, and thus for a valid signature forgery $\mathbf{d}^*$,

$$\mathsf{A}_0\, \mathbf{e}_0 \;=\; \mathsf{A}_0\left(\mathbf{d}_1^* + \mathsf{R}^*\, \mathbf{d}_2^*\right) \;=\; \left[\mathsf{A}_0 \mid \mathsf{A}_0\, \mathsf{R}^*\right] \begin{bmatrix}\mathbf{d}_1^* \\ \mathbf{d}_2^*\end{bmatrix} \;=\; \left[\mathsf{A}_0 \mid \mathsf{C}^*\right] \mathbf{d}^* \;=\; 0 \pmod{q}$$

Next, we show that $\mathbf{e}_0$ is suitably short, which is true since R^* is a sum of $\ell+1$ low-norm matrices R_i with coefficients ± 1, where the summands are all short discrete Gaussian by construction of $\mathsf{R}_0, \ldots, \mathsf{R}_\ell$. Since the matrices $\pm\mathsf{R}_i$ are nearly independent[1] with the same variance $\mathrm{V}\{\pm\mathsf{R}_i\} = \mathrm{V}\{\mathsf{R}_1\}$, we have,

$$\mathrm{V}\{\mathsf{R}^*\} \;=\; \mathrm{V}\{\sum_{i=0}^{\ell} \pm\mathsf{R}_i\} \;\approx\; \sum_{i=0}^{\ell} \mathrm{V}\{\pm\mathsf{R}_i\} \;=\; \sum_{i=0}^{\ell} \mathrm{V}\{\mathsf{R}_i\} \;=\; (\ell+1)\cdot \mathrm{V}\{\mathsf{R}_1\}$$

Since the $\pm\mathsf{R}_i$ closely approximate real normal Gaussian variables, so does R^* and therefore the Gaussian "vanishing tail" inequalities apply. Especially, as they are almost independent discrete Gaussian with center 0 and parameter η, and thus $\mathrm{E}\{\mathsf{R}^*\} \approx \mathrm{E}\{\mathsf{R}_i\} = 0$, we have $\Pr\{\|\pm\mathsf{R}_i\| > \sqrt{m}\,\eta\} = \mathrm{negl}(m)$; and thus,[1]

$$\Pr\{\|\mathsf{R}^*\| > \sqrt{\ell+1}\cdot\sqrt{m}\,\eta\} \;\le\; \Pr\{\|\mathsf{R}_1\| > \sqrt{m}\,\eta\} \;=\; \mathrm{negl}(m)$$

Hence with overwhelming probability $\|\mathbf{e}_0\| \le \beta$ for $\beta = \mathrm{poly}(\ell, n, m) = \mathrm{poly}(\lambda)$, provided we set,[1]

$$\beta \;=\; \left(1 + \sqrt{\ell+1}\,\sqrt{m}\,\eta\right)\sqrt{2\,m}\,\sigma$$

Finally, it remains to show that $\mathbf{e}_0 = \mathbf{d}_1^* + \mathsf{R}_{\mathsf{msg}^*}\, \mathbf{d}_2^* \neq \mathbf{0}$. Suppose for an easy case that $\mathbf{d}_2^* = \mathbf{0}$; then for a valid forgery we must have $\mathbf{d}_1^* \neq \mathbf{0}$ and thus $\mathbf{e}_0 \neq \mathbf{0}$. Suppose on the contrary that $\mathbf{d}_2^* \neq \mathbf{0}$. In that case, $0 \neq \|\mathbf{d}_2^*\| < \sqrt{2\,m}\,\sigma \ll q$; and thus there must be at least one coordinate of $\mathbf{d}_2^*$ that is non-zero modulo q. W.l.o.g., let this coordinate be the last one in $\mathbf{d}_2^*$, and call it y. Let $\mathbf{r}^*$ be the last column of R^*, and let $\mathbf{r}_i$ the last column of R_i for each i. As $\mathsf{R}^* = \sum (-1)^{\mathsf{msg}[i]} \mathsf{R}_i$, we have $\mathbf{r}^* = \sum (-1)^{\mathsf{msg}[i]} \mathbf{r}_i$, where the coefficients ± 1 depend on the message bits. We focus on $\mathbf{r}_1$: the last column of the matrix R_1 associated with the first message bit $\mathsf{msg}[1]$. Let $\mathbf{v} = (-1)^{\mathsf{msg}[1]}\, \mathsf{y}\, \mathbf{r}_1$. The expression of $\mathbf{e}_0$ can be rewritten $\mathbf{e}_0 = \mathsf{y}\, \mathbf{r}^* + \mathbf{e}_0' = \mathbf{v} + \mathbf{e}_0''$, where $\mathbf{v}$ depends on $\mathbf{r}_1$ and $\mathbf{e}_0''$ does not.

The last step is to observe that the only information about $\mathbf{r}_1$ available to $\mathcal{A}$ is contained in the last column of C_1 (with "pollution" $\mathsf{h}_1\, \mathsf{B}_0$, known in the worst case). By leftover hash or a simple pigeonhole principle, there are a very large (exponential in $m - n \log q$) number of admissible and equally likely vectors $\mathbf{r}_1$ that are compatible with the view of $\mathcal{A}$, and in particular more than six of them. Since $\mathcal{A}$ can set the bit $\mathsf{msg}[1]$ in one of two ways, it follows that $\mathcal{A}$ cannot know the value of $\mathbf{v}$ with probability exceeding one third. At most one such value can result in a cancellation of $\mathbf{e}_0$, for if some $\mathbf{v}$ caused all coordinates of $\mathbf{e}_0$ to cancel, then every other $\mathbf{v}$ would fail to do so. We deduce that $\pi_0 = \Pr\{\mathbf{e}_0 \neq 0\} \ge 2/3$. (In fact, we have $\pi_0 > 1 - \exp(-\Omega(m - n\log q)) \to 1$ as $\lambda \to \infty$.) $\square$

[1] Without using near independence, we can show $\Pr\{\|\mathsf{R}^*\| > (\ell+1)\cdot\sqrt{m}\,\eta\} = \mathrm{negl}(m)$, and accordingly set $\beta = \left(1 + (\ell+1)\sqrt{m}\,\eta\right)\sqrt{2\,m}\,\sigma$, which is a factor $\approx \sqrt{\ell}$ worse.

Lemma 27. *For a prime modulus $q = q(\lambda)$ and a number of queries $Q \geq 0$, the simulation completes both the Queries and Forgery phases without aborting, with probability,*

$$\frac{1}{q}\left(1 - \frac{Q}{q}\right) \;\leq\; \Pr\{\mathit{completion}\} \;\leq\; \frac{1}{q}$$

In particular, for $Q \leq q/2$, this probability is $\Pr\{\mathit{completion}\} \in \left[q^{-1}/2,\; q^{-1}\right]$ regardless of the adversary's strategy.

Intuitively, we first observe that *provided* $\mathcal{B}$ does not abort, then the simulation is (almost) perfect in the sense that the view of $\mathcal{A}$ has the same distribution as in an attack against the real scheme (modulo a negligible sampling error owing to the imperfection of $\mathsf{TrapGen}$). In particular, $\mathcal{A}$'s view remains independent of $\mathcal{B}$'s choice of $\mathsf{h}_1, \ldots, \mathsf{h}_\ell$, simply because those values have no counterpart in an actual attack environment.

Now, the adversary can always *assume* that it is facing a simulator instead of a real challenger, and accordingly attempt to derail the simulation. Since the necessity to abort, for a given adversarial strategy, hinges entirely on $\mathcal{B}$'s secret choice of random $\mathsf{h}_1, \ldots, \mathsf{h}_\ell$, it suffices to show that these values remain mostly unlearnable no matter $\mathcal{A}$'s attack strategy.

To show this, we consider a hypothetical unbounded perfect adversary $\mathcal{A}$ and show that, even with perfect Bayesian updating upon each new adaptive query it makes, such adversary is unable to infer enough information about $\mathsf{h}_1, \ldots, \mathsf{h}_\ell$ to affect significantly the success probability of the simulation.

Proof. Consider the ℓ-dimensional space $\mathbb{Z}_q^\ell$, which is the domain of the unknown $(\mathsf{h}_1, \ldots, \mathsf{h}_\ell)$, and recall that $\mathsf{h}_0 = 1$. Denote by $\mathcal{H}_j$ the distribution of $(\mathsf{h}_1, \ldots, \mathsf{h}_\ell)$ over $\mathbb{Z}_q^\ell$ as perceived by the adversary after the first j signature queries have been answered without aborting.

At the start of the attack, since the simulator's selection of $(\mathsf{h}_1, \ldots, \mathsf{h}_\ell)$ is a uniformly random point in $\mathbb{Z}_q^\ell$, the adversary's prior distribution $\mathcal{H}_0$ is necessarily the uniform distribution $\mathcal{U}(\mathbb{Z}_q^\ell)$ over $\mathbb{Z}_q^\ell$. For every query message msg_j that is answered without aborting, $\mathcal{A}$ can prune from the support of $\mathcal{H}$ every point $(\mathsf{h}_1, \ldots, \mathsf{h}_\ell)$ that lies on the "incompatible" hyperplane $\mathsf{h}_{\mathsf{msg}_j} = 0 \pmod q$.

Denote by $\mathbb{V}_j$ the hyperplane thus eliminated after a successful j-th query. Suppose by induction that $\mathcal{H}_{j-1} = \mathcal{U}(\mathbb{W})$, a uniform distribution over some support set $\mathbb{W} \subseteq \mathbb{Z}_q^\ell$. By conditioning $\mathcal{H}_{j-1}$ on the new evidence gained at the j-th query, namely that $(\mathsf{h}_1, \ldots, \mathsf{h}_\ell) \notin \mathbb{V}_j$, one obtains an updated or posterior distribution $\mathcal{H}_j = \mathcal{U}(\mathbb{W} \setminus \mathbb{V}_j)$, which is uniform over the smaller support set given by $\mathbb{W} \setminus \mathbb{V}_j$. By induction on the number of queries, starting from $\mathcal{H}_0 = \mathcal{U}(\mathbb{Z}_q^\ell)$, we deduce that, after the j-th query, $\mathcal{H}_j = \mathcal{U}(\mathbb{Z}_q^\ell \setminus \cup_{i=1}^{j} \mathbb{V}_i)$.

In particular, after all Q allowed queries have been made, the fully updated posterior distribution $\mathcal{H}_Q$ in the view of the adversary is then,

$$\mathcal{H}_Q \;=\; \mathcal{U}(\mathbb{Z}_q^\ell \setminus \cup_{i=1}^{Q} \mathbb{V}_i)$$

In other words, this shows that, in the event that $\mathcal{B}$ was able to answer all the queries, the unknown vector $(\mathsf{h}_1, \ldots, \mathsf{h}_\ell)$ remains equally likely to lie anywhere in all of $\mathbb{Z}_q^\ell$ outside of the Q query-dependent hyperplanes $\mathbb{V}_1, \ldots, \mathbb{V}_Q$. Being the result of perfect Bayesian updating from all available observations, this distribution captures all the information about $(\mathsf{h}_1, \ldots, \mathsf{h}_\ell)$ leaked by $\mathcal{B}$ to $\mathcal{A}$ during the Queries phase.

To complete the argument, consider the hyperplane $\mathbb{V}^* \subset \mathbb{Z}_q^\ell$ defined by the scalar equation $\mathsf{h}_{\mathsf{msg}^*} = 0 \pmod q$ corresponding to the forgery message msg^* chosen by the adversary. By the requirements of what constitutes a valid existential forgery, we know that $\mathsf{msg}^* \neq \mathsf{msg}_j$ and thus $\mathbb{V}^* \neq \mathbb{V}_j$ for all j. (Indeed, the purpose of adding a fixed dummy message bit $\mathsf{msg}[0]$ and setting $\mathsf{h}_0 = 1 \neq 0$ is to ensure that any two distinct messages $\mathsf{msg}_j \neq \mathsf{msg}^* \in \{0,1\}^\ell$ always induce distinct hyperplaces $\mathbb{V}_j \neq \mathbb{V}^* \subset \mathbb{Z}_q^\ell$.)

Since $\mathbb{V}^*$ and $\mathbb{V}_j$ are distinct affine subspaces of dimension $\ell - 1$ in $\mathbb{Z}_q^\ell$, we have $\left|\mathbb{V}^* \cap \mathbb{V}_j\right| \leq q^{\ell-2}$ whereas $\left|\mathbb{V}^*\right| = \left|\mathbb{V}_j\right| = q^{\ell-1}$ and of course $\left|\mathbb{Z}_q^\ell\right| = q^\ell$. Consequently, $\mathbb{V}^*$ and $\mathbb{V}_j$ have at most a fraction $1/q$ of their points in common, and more specifically $\left|\mathbb{V}^* \setminus \mathbb{V}_j\right| \geq (1 - q^{-1}) \left|\mathbb{V}^*\right| = (1 - q^{-1})\, q^{-1} \left|\mathbb{Z}_q^\ell\right|$ for all j.

Considering the event *completion* $= \wedge_{i=1}^{Q} \big\{(\mathsf{h}_1, \ldots, \mathsf{h}_\ell) \in (\mathbb{V}^* \setminus \mathbb{V}_i)\big\}$ and invoking the union bound on this conjunction, we thus establish a lower bound,

$$\begin{aligned} \Pr\big\{\mathit{completion}\big\} &= \Pr\big\{(\mathsf{h}_1, \ldots, \mathsf{h}_\ell) \in (\mathbb{V}^* \setminus \cup_{i=1}^{Q} \mathbb{V}_i)\big\} \\ &\geq \big(1 - q^{-1}\, Q\big) \Pr\big\{(\mathsf{h}_1, \ldots, \mathsf{h}_\ell) \in \mathbb{V}^*\big\} \;=\; \big(1 - q^{-1}\, Q\big)\, q^{-1} \end{aligned}$$

Conversely, we can trivially establish an upper bound,

$$\begin{aligned} \Pr\big\{\mathit{completion}\big\} &= \Pr\big\{(\mathsf{h}_1, \ldots, \mathsf{h}_\ell) \in (\mathbb{V}^* \setminus \cup_{i=1}^{Q} \mathbb{V}_i)\big\} \\ &\leq \Pr\big\{(\mathsf{h}_1, \ldots, \mathsf{h}_\ell) \in \mathbb{V}^*\big\} \;=\; \left|\mathbb{V}^*\right| \,/\, \left|\mathbb{Z}_q^\ell\right| = q^{-1} \end{aligned}$$

In both cases the probability is over the simulator's initial choice of $\mathsf{h}_1, \ldots, \mathsf{h}_\ell$.

We have shown that the probability of *completion* without aborting is bounded in the narrow range $\big[(1 - q^{-1}\, Q)\, q^{-1},\ q^{-1}\big]$, regardless of the adversary's actions. The lemma follows. $\square$

3.4 Lattice Parameters

It is not so obvious to see that the various parameters can be instantiated in a way that satisfies the flurry of constraints and inequalities evoked in § 2 and § 3.3. This is necessary for us, later, to prove the security of the signature from a polynomial average-case SIS that reduces to a worst-case lattice hardness assumption.

Example 28. To ensure that hard lattices with good short bases can be generated (i.e., $m \geq 6\, n\, \log q$), that our flavor of SIS has a worst-case lattice reduction (i.e., $q \geq \beta \cdot \omega(\sqrt{n \log n})$), that the two-sided trapdoors can operate smooothly (i.e., σ sufficiently large), that vectors sampled using a trapdoor are difficult SIS solutions (i.e., $\beta \geq \sqrt{2\, \ell}\, m\, \eta\, \sigma$), etc., in function of a security parameter λ: we may choose a function $\omega(\sqrt{\log m})$, a constant $\delta_1 > 0$, and a threshold $\lambda_0 \gg 0$; and $\forall \lambda > \lambda_0$ may set:

$$\begin{aligned} n &= \lambda \\ m &= n^{1+\delta_1} \\ \eta &= \omega(\sqrt{\log m}) \\ L &= \sqrt{m}\,\omega(\sqrt{\log m}) \\ \sigma &= \sqrt{\ell}\,m^{3/2}\,\omega(\sqrt{\log m})^4 \\ \beta &= \sqrt{2}\,\ell\,m^{5/2}\,\omega(\sqrt{\log m})^5 \\ q &= \sqrt{2}\,\ell\,m^3\,\omega(\sqrt{\log m})^6 \end{aligned}$$

One must however keep in mind that the security reduction given in Theorem 25 holds only if $q \geq 2\,Q$, so it may be necessary to increase q and the other parameters beyond the baseline values listed above. We avoid this in § 3.5.

3.5 Refined Simulation Framework

In the full paper, we give a refined analysis of such schemes that lets us retain the baseline q even for very large Q. The idea is to replace the random scalars $\mathsf{h}_i \in \mathbb{Z}_q$ by block-diagonal matrices $\mathsf{H}_i \in \mathbb{Z}_q^{n\times n}$ consisting of a repeated random submatrix drawn from a *full-rank difference* group $\mathcal{G} \subset \mathbb{Z}_q^{k\times k}$ for a special $k|n$, where any difference $G_1 - G_2 \in \mathcal{G}$ is either zero or an invertible matrix in $\mathbb{Z}_q^{k\times k}$. Visually, a random input k-vector $\in \mathbb{Z}_q^k$ is mapped to a random matrix H_i using an encoding map μ built from an FRD encoding φ, according to this picture,

$$\mu \;:\; \mathbb{Z}_q^k \to \mathbb{Z}_q^{n\times n} \;:\; \mathbf{v} \mapsto \begin{pmatrix} \boxed{\varphi(\mathbf{v})} & & 0 \\ & \underbrace{\boxed{\varphi(\mathbf{v})}}_{k}\}^{k} & \\ 0 & & \ddots \;\boxed{\varphi(\mathbf{v})} \end{pmatrix}$$

Full-rank difference (FRD) families in $\mathbb{Z}_q^{n\times n}$ were used as a plentiful IBE encoding [1] able to represent *as many as possible* distinct identities, up to q^n.

Here, FRD families serve a very different purpose, internal to the simulator. They let us swap the mixing scalars h_i for *uniform* matrices H_i in a custom group, whose size q^k can be tuned *just right* to the query allowance Q without also raising the modulus q. The benefit: smaller moduli make signatures smaller and faster, and security tighter. Remarkably, except for the relaxation on q, the actual scheme is unchanged. The theorem below is proven in the full paper.

Theorem 29. *If there exists a probabilistic algorithm $\mathcal{A}$ that creates an existential signature forgery, in time τ, with probability ϵ, making Q adaptive chosen-message queries, then there exists a probabilistic algorithm $\mathcal{B}$ that solves the SIS problem of Theorem 25 in time $\tau' \approx \tau$ with probability $\epsilon' \geq \epsilon/(6\,q\,Q)$.*

Since both versions of the framework apply to the same scheme, we can pick q obliviously of Q, and invoke Theorem 25 if $Q \leq q/2$ or Theorem 29 if $Q \gg q$.

Acknowledgments The author thanks Shweta Agrawal, Dan Boneh, Ronald Cramer, David Freeman, and anonymous referees for valuable insights.

References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, 2010.
2. S. Agrawal, D. Boneh, and X. Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. Manuscript, 2010.
3. S. Agrawal and X. Boyen. Identity-based encryption from lattices in the standard model. Manuscript, 2009. `http://www.cs.stanford.edu/~xb/ab09/`.
4. D. Aharonov and O. Regev. Lattice problems in NP $\cap$ coNP. *Journal of the ACM*, 52(5):749–65, 2005.
5. M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, 1996.
6. M. Ajtai. Generating hard instances of the short basis problem. In *ICALP*, 1999.
7. J. Alwen and C. Peikert. Generating shorter bases for hard random lattices. In *STACS*, 2009.
8. D. Boneh and X. Boyen. Efficient selective-ID secure identity based encryption without random oracles. In *EUROCRYPT*, 2004.
9. D. Boneh and X. Boyen. Efficient lattice (H)IBE in the standard model from the BB-1 framework. Manuscript, 2009. Slides at `http://rump2009.cr.yp.to/`.
10. R. Canetti, S. Halevi, and J. Katz. A forward-secure public-key encryption scheme. In *EUROCRYPT*, 2003.
11. D. Cash, D. Hofheinz, and E. Kiltz. How to delegate a lattice basis. Cryptology ePrint Archive, Report 2009/351, 2009. `http://eprint.iacr.org/`.
12. N. Gama and P. Q. Nguyen. Predicting lattice reduction. In *EUROCRYPT*, 2008.
13. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
14. S. Hohenberger and B. Waters. Short and stateless signatures from the RSA assumption. In *CRYPTO*, 2009.
15. L. Lovász. *An Algorthmic Theory of Numbers, Graphs and Convexity*. SIAM, 1986.
16. V. Lyubashevsky and D. Micciancio. Asymptotically efficient lattice-based digital signatures. In *TCC*, 2008.
17. D. Micciancio and S. Goldwasser. Complexity of lattice problems: a cryptographic perspective. *Kluwer Series on Engineering and Computer Science*, 2002.
18. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. In *FOCS*, 2004.
19. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM Journal of Computing*, 37(1):267–302, 2007.
20. D. Micciancio and O. Regev. Lattice-based cryptography. In D. J. Bernstein and J. Buchmann, editors, *Post-quantum Cryptography*. Springer, 2008.
21. C. Peikert. Bonsai trees (or, arboriculture in lattice-based cryptography). Cryptology ePrint Archive, Report 2009/359, 2009. `http://eprint.iacr.org/`.
22. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, 2006.
23. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.
24. B. Waters. Efficient identity-based encryption without random oracles. In *EUROCRYPT*, 2005.